

~~METHOD AND APPARATUS FOR SAVING AND RESTORING
PROCESSOR REGISTER VALUES AND ALLOCATING AND
DEALLOCATING STACK MEMORY~~

Inventors: Kevin D. Kissell
Hartvig W. J. Ekner

CROSS REFERENCE TO RELATED APPLICATION

This application contains subject matter related to commonly owned, co-pending U.S. Patent Application Ser. No. 09/836,541, filed April 14, 2001, entitled "Mapping System and Method for Instruction Set Processing," which is herein incorporated by reference in its entirety. ✓

BACKGROUND OF THE INVENTION

Field of the Invention

[0001] The invention relates generally to a data processing system. More particularly, it relates to saving and restoring processor register values and allocating and deallocating stack memory.

Background Art

[0002] Two well-known operations performed by computer systems are the storing and retrieving of items on a stack. Stackable items include general purpose register contents; e.g., data and addresses. These operations (also referred to as "push" and "pop" operations) are typically used to facilitate the entry to and exit from subroutines. That portion of a stack created for a particular subroutine is referred to as a "stack frame." In programmable devices (such as microprocessors), dedicated instructions may be used to carry out these operations.

[0003] It is desired to enhance the utility of stack storing and/or retrieving operations by providing additional functionality associated therewith. Such functionality, when added to instructions for carrying out stack operations, make it possible to write more compact application programs since such instructions encode multiple functions.

BRIEF SUMMARY OF THE INVENTION

[0004] The present invention provides methods and means for saving and restoring processor registers and allocating and deallocating a stack frame. In one embodiment, a first field of a save instruction encodes whether a value in a register of a processor is saved as an argument value or a static value. A second field of the save instruction encodes a size of a stack frame created during execution of the save instruction. An argument value is saved in a calling program's stack frame. A static value is saved in a called program's stack frame. A restore instruction is used to restore a static value and deallocate the stack frame. The save and restore instructions may be executed using any programmable device, including a single instruction set architecture processor or a multi-instruction set architecture processor. The functionality of such instructions may be achieved through software, hardware or a combination of both.

[0005] In another embodiment, a 16-bit instruction according to the invention comprises at least five fields. These five fields are an instruction opcode field, a 1-bit return address register field, a 1-bit first static register field, a 1-bit second static register field, and a 4-bit frame-size field. This instruction can be executed as a single 16-bit instruction or executed in combination with a 16-bit instruction extension. An instruction extension comprises at least four fields. These four fields are an extend instruction opcode field, a 3-bit additional static registers field, a second 4-bit frame-size field, and a 4-bit arguments register field. The 3-bit additional static registers field allows the values in up to seven addition registers to be saved and restored as static values.

[0006] Features of the invention allow the invention to be implemented, for example, as a method for encoding an instruction, as a processor core, as a mapper, as a decoder, and/or as a computer program.

[0007] Further features and advantages of the present invention, as well as the structure and operation of various embodiments of the present invention, are described in detail below with reference to the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS/FIGURES

[0008] The features and advantages of the present invention will become more apparent from the detailed description set forth below when taken in conjunction with the drawings in which like reference numbers indicate identical or functionally similar elements.

[0009] FIG. 1 is a block diagram illustrating a single-instruction set architecture processor system according to an embodiment of the invention.

[0010] FIG. 2 illustrates an exemplary register bank according to an embodiment of the invention.

[0011] FIG. 3 illustrates an exemplary program stack in memory.

[0012] FIG. 4 illustrates two exemplary instructions for allocating memory and saving registers of a processor core on entry to a subroutine according to the invention.

[0013] FIGs. 5A-^{5B}~~B~~ illustrate a flowchart of a method for saving registers and allocating memory on entry to a subroutine according to an embodiment of the invention.

[0014] FIG. 6 illustrates an example computer program pseudo-code for a save instruction according to the invention.

[0015] FIG. 7 illustrates an example encoding of an argument registers field of an instruction according to the invention.

[0016] FIGs. 8A-8B illustrate an example computer program pseudo-code for an extended save instruction according to the invention.

103K
7-15-04

[0017] FIG. 9 illustrates two exemplary instructions for deallocating memory and restoring registers of a processor core before exit from a subroutine according to the invention. 10B

[0018] FIGs. 10A-~~B~~^A illustrate a flowchart of a method for restoring registers and deallocating memory before exit from a subroutine according to an embodiment of the invention.

KSK
7/15/04

[0019] FIG. 11 illustrates an example computer program pseudo-code for a restore instruction according to the invention.

[0020] FIGs. 12A-12B illustrate an example computer program pseudo-code for an extended restore instruction according to the invention.

[0021] FIGs. 13^{14, 15, 16, 17, and}~~18~~^A illustrate the operation of example save and restore instructions encoded according to embodiments of the invention.

"

DETAILED DESCRIPTION OF THE INVENTION

[0022] The present invention provides methods and means for saving and restoring a processor register value and allocating and deallocating stack memory. Method embodiments of the invention are encoded in instructions that can be efficiently stored as a part of a computer program. Apparatus embodiments of the invention execute instructions according to the invention. As described herein, instructions according to the invention can be executed using either a single instruction set architecture processor or a multi-instruction set architecture processor.

[0023] In an embodiment of the invention, a first field of a save instruction encodes whether a value in a register of a processor is saved as an argument value or a static value. A second field of the save instruction encodes a size of a stack frame created during execution of the save instruction. An argument value is saved in a calling program's stack frame. A static value is saved in a called program's stack frame. A restore instruction is used to restore a static value and deallocate the stack frame.

[0024] The invention is now described with reference to FIGs. 1-18.

Example Apparatus/System Embodiments of the invention

[0025] FIG. 1 depicts an example environment for implementation of the present invention. Specifically, FIG. 1 depicts a processor system 100 for implementing the present invention.

[0026] System 100 comprises a processor 101, a cache 102, a cache controller 110, a memory 104, a memory management unit 160, and a processor-memory bus 105. A bus interface 103 couples processor-memory bus 105 to an input/output (I/O) bus 106 and a plurality of I/O controllers such as, for example, I/O controllers 107, 108, and 109. Cache controller 110 has a mapper 120. Processor 101 has a decoder 152 and a register bank 156.

[0027] FIG. 2 illustrates an example register bank 156 having thirty-two registers labeled R0-R31. In other embodiments, register bank 156 may have more or less than thirty-two registers. Certain of the registers of register bank 156 are assigned labels. Registers R4, R5, R6, and R7 are shown labeled as argument registers A0, A1, A2, and A3, respectively. Registers R16 and R17 are labeled as static registers S0 and S1, respectively. A static register is a processor register which, by software convention, contains a data value that must be preserved across subroutine calls. A subroutine that respects a static register convention must exit with all static registers in the same state that they were in at subroutine entry. This can be achieved either by never modifying the static registers, or by saving the static value before any modified use of the register and then restoring the value before exit from the subroutine. Register R29 is labeled as stack pointer register SP. Register R31 is labeled as return address register RA. These labels are illustrative, and not intended to limit the invention. The labels are used below to further describe the invention. Registers other than the registers labeled in FIG. 2 can be used to implement the invention, as will be understood by a person skilled in the relevant art given the description herein.

[0028] FIG. 3 illustrates a portion of memory 104 in greater detail. Memory 104 comprises a user memory 302 and an operating system memory (not shown). User memory 302 comprises five regions of interest. These regions are a program

code region 304, a data region 306A, a heap region 306B, a stack region 308, and an unused or available space region 310. As can be seen in FIG. 3, stack region 308 typically comprise a region 312 for passing argument values to a subroutine via memory 104, and a reserved space for storing values passed to a subroutine in a register. A stack pointer 316 indicates the address of the next available memory location in region 310 of memory 104.

[0029] The program code region 304 is used to store the instructions of a program being implemented by system 100. Program instructions to be executed by the processor of system 100 must be read into cache memory 102 from main memory 104.

[0030] Heap 306B is used for globally meaningful or long-lived dynamic data, while stack 308 is used for locally significant and short-lived dynamic data. Allocation of heap memory 306B involves software manipulation of data structures in memory that control its explicit allocation. Allocation of dynamic data on stack 308 is done simply by updating the stack pointer 316 and remembering the stack pointer offset of each allocated data item. As shown in FIG. 3, stack 308 grows from higher memory addresses to lower memory address. In other embodiments of the invention, stack 308 can grow from lower memory addresses to higher memory addresses. The data region 306A is used to store static data.

[0031] Further description of the elements of system 100 can be found in any general computer design book. For example, see "Computer Organization & Design: the Hardware/Software Interface," by David A. Patterson and John L. Hennessy, 2nd ed., Morgan Kaufmann Publishers Inc., San Francisco, California, 1998, which is incorporated herein by reference in its entirety for all purposes.

[0032] The present invention is now described with reference to the features of system 100 which provides an example environment for implementation of the invention. It will be apparent to a person skilled in the art, however, that the invention can be implemented in other processor environments. For example, the invention can be implemented in a dual instruction set processor environment such as that described in commonly owned, co-pending U.S. Pat. Appl. No.

09/836,541, filed April 14, 2001, and titled "Mapping System and Method for Instruction Set Processing," which is incorporated herein by reference as if reproduced in full below.

Example Method Embodiments of the Invention

[0033] As described herein, the steps of method embodiments of the invention can be encoded in instructions that are stored as a part of a computer program. These methods are then implemented using a processor (e.g., processor 101). A processor according to the invention can decode the instructions described herein and perform the steps of the method embodiments.

[0034] FIG. 4 illustrates two exemplary 16-bit instructions 402 and 452 that encode method embodiments of the invention for saving registers and allocating stack memory upon entry to a subroutine. Instruction 402 can be executed as a single, stand-alone 16-bit instruction. In addition, instruction 402 can be executed in combination with instruction extension 452 as a 32-bit extended version of the 16-bit instruction (i.e., instruction 475).

[0035] Instruction 402 comprises five fields. These five fields are a 9-bit save instruction opcode field 404, a 1-bit return address register (RA) field 406, a 1-bit first static register (S0) field 408, a 1-bit second static register (S1) field 410, and a 4-bit frame-size field 412. As would be known to a person skilled in the relevant art, a return address register is used to store an address of an instruction that is to be executed upon exit from a subroutine, and a static register is used to store a value of a variable that must be preserved by a program before calling a subroutine. The purpose of each of these fields is described below with regard to the steps of method 500 and pseudo-code 650.

[0036] Instruction extension 452 comprises four fields. These four fields are a 5-bit extend instruction opcode field 454, a 3-bit additional static registers field 456, a 4-bit frame-size field 458, and a 4-bit argument registers field 460. As would be known to a person skilled in the relevant art, an argument register is used by a calling program to pass a value to a subroutine. The value in the

argument register may or may not need to be preserved by the subroutine, depending on software convention. The purpose of each of these fields is described below with regard to the steps of method 500 and pseudo-code 800.

[0037] A method 500 for saving registers and allocating memory on entry to a subroutine according to the invention will now be described. FIGs. 5A and 5B illustrate a flowchart of the steps of method 500. As illustrated in the flowchart, method 500 comprises a plurality of steps divided up into three stages 501, 509, and 515. Stage 501 comprises steps 502, 504, 506, and 508, and is performed prior to or during development and/or implementation of a processor design. As will be understood by a person skilled in the relevant art(s), once these steps have been performed, they need not be repeated. Steps 502-508 are an inherent part of a processor, a mapper, a decoder and/or an instruction according to the invention. Stage 509 comprises steps 510, 512, and 514, and an encoding step implicit in step 522. These steps of method 500 are typically performed using a compiler program. A compiler program takes source code and uses the source code to determine a desired encoding for instructions according to the invention. Lastly, stage 515 comprises steps 516, 518, 520, and 522. These steps are performed by an apparatus or system according to the invention executing an instruction encoded according to the invention. As described herein, it is not necessary to perform each step of method 500 every time the method is implemented.

[0038] The steps of method 500 can be encoded by instruction 402 and/or instruction 475 for implementation by a processor. In the description that follows, method 500 is described in detail with regard to instruction 402, register bank 156, stack region 308 of memory 104, and pseudo-code 650 (shown in FIG. 6). Method 500 begins at a step 502.

[0039] In step 502, at least one general purpose register of register bank 156 is specified as an argument register. This is to enable subsequent steps of method 500 to be encoded using the fields of instructions 402 and 475. The number of argument registers specified will typically depend on the number of general purpose registers available in a processor core and/or the number of registers

intended to be visible when using a particular instruction set. In an embodiment of the invention, registers R4, R5, R6, and R7 are specified as argument registers A0, A1, A2, and A3, respectively, as shown in FIG. 2. In other embodiments, more or less than four argument registers can be specified in step 502.

[0040] Specifying a general purpose register of register bank 156 in step 502 as an argument register is not intended to alter the nature of the general purpose register with regard to other instructions of an instruction set. In fact, a feature of the invention allows any specified argument register according to the invention to be treated as a general purpose register by instructions other than the instructions described herein, and thus maintain all the functionality of a general purpose register with regard to other instructions. Furthermore, this feature of the invention also applies to general purpose registers specified, for example, as a static register (e.g., S0 and S1), a stack pointer register (SP), and/or a return address register (RA) in accordance with method 500.

[0041] In step 504, at least one register of register bank 156 is specified as a static register. In the embodiment of FIG. 2, two registers R16 and R17 are specified as static registers S0 and S1, respectively. In other embodiments, more or less than two registers may be specified.

[0042] In step 506, a register of register bank 156 is specified as a stack pointer register (SP). The stack pointer register typically holds the address value of the next available memory location in memory 104. FIG. 3 illustrates a stack pointer 316. In the embodiment shown in FIG. 2, register R29 of register bank 156 is specified as the stack pointer register.

[0043] In step 508, a register of register bank 156 is specified as a return address register (RA). The return address register typically holds the address value of the next instruction to be executed upon exit from a subroutine. In the embodiment shown in FIG. 2, register R31 is specified as the return address register.

[0044] As would be understood by a person skilled in the relevant art, the purpose of specifying various registers of register bank 156 as either an argument register, a static register, a stack pointer register, or a return address register is to enable programs to be developed that can compile and link other program code

written in a high level programming language such as, for example, C, C++, Pascal, and/or Fortran to run on a processor according to the invention. The number of registers specified in step 502 as argument registers and the number of registers specified in step 504 as static registers, however, may not be optimal for all applications. Thus, as described below, the invention enables a compiler program to encode in an instruction whether any or all of the argument registers specified in step 502 should be handled as static registers in order to enhance a particular feature of an application program. Method 500 enables a compiler program to encode in an instruction, for example, whether a value in a specified argument register is to be saved at one of two locations in a stack memory. How this is accomplished is described with regard to the following steps of method 500.

[0045] In step 510, fields 408 and 410 of instruction 402 are encoded to indicate whether static registers R16 and R17 are to be saved to region 310 of memory 104 upon entry to a subroutine. If a called subroutine will use a register that has been specified as a static register in step 504, the value of the specified static register is saved on stack 308 in order to preserve the value for the program that called the subroutine. For example, if a call subroutine will use register R16 of register bank 156, the value of register R16 is saved on stack 308 upon entry to the subroutine. A compiler program can ensure that the value of register R16 is saved on stack 308 by encoding a value in field 408 of instruction 402 (e.g., by setting bit 5 of instruction 402) that instructs a processor to save the value of register R16 on stack 308. In an embodiment, if the compiler clears bit 5 of instruction 402, the value of register R16 will not be saved on stack 308 when instruction 402 is executed. Similarly, whether a value in register R17 is saved on stack 308 when instruction 402 is executed by a processor is determined by the state of bit 4 of instruction 402. In an embodiment, if bit 4 is set, for example, the value of register R17 will be saved on stack 308 when instruction 402 is executed. If bit 4 is cleared, the value of register R17 will not be saved on stack 308 when instruction 402 is executed.

[0046] For the embodiment shown in FIG. 4, two 1-bit fields 408 and 410 are used to encode whether the values of two specified static registers are saved on

stack 308. In other embodiments, a single multi-bit field can be used to encode whether values of specified static registers are saved. How to do this will be understood by a person skilled in the relevant art given the description herein.

[0047] Step 512 of method 500 is not performed by the method embodiment encoded by instruction 402. For the method embodiment encoded by instruction 402, it is assumed that each of the registers specified in step 502 are to be treated as argument registers by a called subroutine. This is not the case, however, for other method embodiments of the invention encoded, for example, using instruction 475. Step 512 is further described below with regard to instruction 475.

[0048] In step 514, field 406 of instruction 402 is encoded to indicate whether a value in return address register R31 is to be saved to region 310 of memory 104 upon entry to a subroutine. In an embodiment, for example, the value of register R31 will be saved on stack 308 when instruction 402 is executed if bit 6 of instruction 402 is set. If bit 6 is cleared, the value of register R31 will not be saved on stack 308 when instruction 402 is executed.

[0049] Step 516 of method 500 is also not performed by the method embodiment encoded by instruction 402. As described with regard to step 512, for the method encoded by instruction 402, it is assumed that each of the registers specified in step 502 is to be treated as an argument register by a called subroutine. Thus, none of the argument registers R4, R5, R6, or R7 is saved as a static upon entry to a subroutine.

[0050] In step 518, a value in register R16 is either saved or not saved on stack 308 when instruction 402 is executed. Whether the value in register R16 is saved is based on the encoding of field 408. If the value in register R16 is to be saved on stack 308, the address of the memory location where the value is saved may be calculated any number of ways, including for example the way shown in the pseudo-code of FIGs. 6 and/or 8A and 8B. In an embodiment, the value in register R16 is saved on stack 308 only if bit 5 of instruction 402 is set. Otherwise, the value of register R16 will not be saved when instruction 402 is executed. Similarly, in step 518, a value in register R17 will be saved on stack

308 when instruction 402 is executed only if bit 4 of instruction 402 is set. Otherwise, the value in register R17 will not be saved when instruction 402 is executed. If the value in register R17 is to be saved on stack 308, the address of the memory location where the value is saved may be calculated any number of ways, including for example the way shown in the pseudo-code of FIGs. 6 and/or 8A and 8B.


[0051] In step 520, a value in register R31 is either saved or not saved on stack 308 when instruction 402 is executed. Whether the value in register R31 is saved is based on the encoding of field 406. If the value in register R31 is to be saved on stack 308, the address of the memory location where the value is saved may be calculated any number of ways, including for example the way shown in the pseudo-code of FIGs. 6 and/or 8A and 8B. In an embodiment, the value in register R31 is saved on stack 308 only if bit 6 of instruction 402 is set. Otherwise, the value of register R31 will not be saved when instruction 402 is executed.

[0052] Lastly, in step 522, the value in the specified stack pointer register (R29) is adjusted based on a value encoded in field 412 of instruction 402. Adjusting the stack pointer allows a stack frame to be set up and memory to be allocated between the calling program and the called subroutine. In an embodiment, the value in the specified stack pointer register is adjusted by subtracting 8 times the value encoded in field 412 of instruction 402. If the value encoded in field 412 is zero, the value in the stack pointer register is adjusted by subtracting 128. In other embodiments, the value in the stack pointer register is adjusted by different amounts.

[0053] FIG. 6 illustrates pseudo-code 650 according to the invention. Pseudo-code 650 describes to a person skilled in the relevant art how a processor according to the invention operates to implement the steps of method 500 encoded in an instruction 402.

[0054] The steps of pseudo-code 650 will now be described. The steps of pseudo-code 650 are described with reference to register bank 156, stack 1300, and the example instruction 402A shown below and in FIG. 13. Stack 1300 is

shown having memory locations M0-M14. As can be seen by examining the example instruction, the encodings of the instruction fields are: "1" for return address field 406; "1" for static register field 408; "1" for static register field 410; and "0100" (i.e., 4) for frame-size field 412. The encoding for the save instruction opcode field 404 is shown as "XXXXXXXXXX" to indicate that the encoding is a processor specific value that is unimportant for understanding the steps of pseudo-code 650.



[0055] Implementation of pseudo-code 650 starts by storing the value of register R29 (the specified stack pointer) to a temporary variable. The temporary variables in pseudo-code do not represent program variables in memory or registers, but rather temporary values of internal state within a processor. At the start of pseudo-code 650, the stack pointer is assumed to have a value equal to that of stack memory location M12. After storing the stack pointer value, the value of return address field 406 is examined. In the case of the example instruction above, the encoded value is 1. Thus, the value of the temporary variable is reduced by four and the value in register R31 (the return address) is saved at memory location M11. Next, the value of static field 410 is examined. In the case of the example instruction above, the encoded value is 1. Thus, the value of the temporary variable is again reduced by four and the value in register R17 (register S1) is saved at memory location M10. Next, the value of static field 408 is examined. In the case of the example instruction above, the encoded value is 1. Thus, the value of the temporary variable is again reduced by four and the value in register R16 (register S0) is saved at memory location M9. Lastly, the value of frame-size field 412 is examined. In the case of the example instruction above, the encoded value is 4. As indicated by pseudo-code 650, the value in the specified stack pointer register (R29) is adjusted by shifting the value in frame-size field 412 left three bits, padding the new value with zeros, and then subtracting the new value from the value in register R29. At the end of this

operation, the stack pointer points at a memory location M4, thereby defining a stack frame from M11 to M4.

[0056] As will be understood by a person skilled in the relevant art, the stack pointer is adjusted in the above example to reserve space, for example, for storing values of the specified argument registers (i.e., R4, R5, R6, and R7) in stack 1300. As shown in FIG. 13, memory locations M4, M5, M6, and M7 can be used to save the values of registers R7, R6, R5, and R4, respectively, if necessary. Additional space, below memory location M4, can be reserved by increasing the value encoded in field 412 of instruction 402A.

[0057] Method 500 will now be described in detail with regard to instruction 475, register bank 156, stack 1400 (shown in FIG. 14), and pseudo-code 800 (shown in FIGs. 8A and 8B) in order to point out differences between instruction 475 and instruction 402.

[0058] As will be understood by a person skilled in the relevant art given the description herein, steps 502, 504, 506, and 508 do not change when instruction 475 is used to encode method 500 rather than instruction 402. Thus, the description for these steps is not repeated here.

[0059] In step 510, fields 408 and 410 of instruction 475 are encoded to indicate whether specified static registers R16 and R17 are to be saved on stack 1400 upon entry to a subroutine. If a called subroutine will use a register that has been specified as a static register in step 504, the value of the specified static register is save on stack 1400 in order to preserve the value for the program that called the subroutine. As described below, in an embodiment the value of register R16 is saved at memory location M5 during execution of instruction 475 by setting bit 5 of instruction 475 to a value of 1. If bit 5 of instruction 475 is 0, the value of register R16 will not be saved on stack 1400 when instruction 475 is executed. Similarly, whether a value in register R17 is saved at a memory location on stack 1400 when instruction 475 is executed by a processor is determined by the state if bit 4 of instruction 475. If bit 4 is set, the value of register R17 will be saved, for example in memory location M6 when instruction 402 is executed. If bit 4 is

cleared, the value of register R17 will not be saved when instruction 402 is executed.

[0060] In step 512 of method 500, a binary value (e.g., 4-bit binary value) is encoded in argument registers field 460 (*aregs*) of instruction extension 452 to indicate which registers specified as argument registers in step 502 are to be treated and saved during execution of instruction 475 as static registers. FIG. 7 illustrates a 4-bit binary encoding used in an embodiment of the invention. Other encodings, however, can also be used in accordance with method 500.

[0061] In step 514, field 406 of instruction 475 is encoded to indicate whether a value in specified return address register R31 is to be saved, for example, at a memory location M11 upon entry to a subroutine. In an embodiment, for example, the value of register R31 will be saved when instruction 475 is executed in memory location M11 if bit 6 of instruction 475 is set. If bit 6 is cleared, the value of register R31 will not be saved when instruction 475 is executed.

[0062] In step 516, none, one, two, three, or four of the registers R4 (A0), R5 (A1), R6 (A2), or R7 (A3) specified as argument registers in step 502 are saved on stack 1400 during execution of instruction 475. Which, if any, specified argument registers are saved depends on the value encoded in field 460 of instruction 475. FIG. 7 shows which specified argument registers are treated as argument registers and which specified argument registers are treated as static registers for the 4-bit binary encoding illustrated in FIG. 7. If one or more values in registers R4, R5, R6, and R7 are to be saved on stack 1400, the addresses of the memory locations where the values are saved may be calculated any number of ways, including for example the way shown in the pseudo-code of FIGs. 8A and 8B.

[0063] In step 518, a value in register R16 (S0) is either saved or not saved on stack 1400 when instruction 475 is executed. Whether the value in register R16 is saved is based on the encoding of field 408. If the value in register R16 is to be saved on stack 1400, the address of the memory location where the value is saved may be calculated any number of ways, including for example the way shown in the pseudo-code of FIGs. 6 and/or 8A and 8B. In an embodiment, the

value in register R16 is saved on stack 1400 only if bit 5 of instruction 475 is set. Otherwise, the value of register R16 will not be saved when instruction 475 is executed. Similarly, in step 518, a value in register R17 (S1) will be saved on stack 1400 when instruction 475 is executed only if bit 4 of instruction 475 is set. Otherwise, the value in register R17 will not be saved when instruction 475 is executed. If the value in register R17 is to be saved on stack 1400, the address of the memory location where the value is saved may be calculated any number of ways, including for example the way shown in the pseudo-code of FIGs. 6 and/or 8A and 8B.


[0064] In step 520, a value in register R31 (RA) is either saved or not saved on stack 1400 when instruction 475 is executed. Whether the value in register R31 is saved is based on the encoding of field 406. If the value in register R31 is to be saved on stack 1400, the address of the memory location where the value is saved may be calculated any number of ways, including for example the way shown in the pseudo-code of FIGs. 6 and/or 8A and 8B. In an embodiment, the value in register R31 is saved on stack 1400 only if bit 6 of instruction 475 is set. Otherwise, the value of register R31 will not be saved when instruction 475 is executed.

[0065] In step 522, the value in register R29 (SP) is adjusted based on a value encoded in fields 458 and 412 of instruction 475. Adjusting the stack pointer value in register R29 allows a stack frame to be set up. In an embodiment, the value in register R29 is adjusted by subtracting 8 times the value encoded in fields 412 and 458 of instruction 475. The 4-bits of field 458 and the 4-bits of field 412 are concatenated to form an 8-bit frame-size value. In other embodiments, the value in register R29 is adjusted by different amounts.

[0066] Field 456 of instruction 475 is used to encode whether additional registers of register bank 156 are to be treated as static registers. For example, in an embodiment, registers R18, R19, R20, R21, R22, R23, and R 30 can be treated as additional static registers. This feature of the invention is further described below with reference to pseudo-code 800.

[0067] FIGs. 8A and 8B illustrate pseudo-code 800 according to the invention. Pseudo-code 800 describes to a person skilled in the relevant art how a processor according to the invention implements the steps of method 500 encoded in an instruction 475.

[0068] The steps of pseudo-code 800 will now be described. The steps of pseudo-code 800 are described with reference to register bank 156, stack 1400, and example instruction 475A shown below and in FIG. 14. As can be seen by examining the example instruction, the encodings of the instruction fields are: "1" for return address field 406; "1" for static register field 408; "1" for static register field 410; "00010100" (i.e., 20) for the concatenated frame-size fields 458 and 412; "100" for the additional static registers field 456; and "1010" for the argument registers field 460. The encodings for the save instruction opcode field 404 and the extend instruction opcode field 454 are shown as "XXXXXXXXXX" and "XXXXXX," respectfully, to indicate that these encodings are processor specific values that are unimportant for understanding the steps of pseudo-code 800.



[0069] Implementation of pseudo-code 800 starts by storing the value of register R29 (the stack pointer) to two temporary variables. At the start of pseudo-code 800, the stack pointer is assumed to be pointing at memory location M12.

[0070] After saving the value of the stack pointer in two temporary variables, the number of argument registers is determined. In the example instruction above, the number encoded in field 460 is "1010." By looking at either FIG. 7 or pseudo-code 800, it can be determined that the values in registers R4 and R5 are to be

treaded and stored as argument values (e.g., they are to be stored as values associated with storage already allocated on the stack). In accordance with the method of the invention described by pseudo-code 800, the values of argument registers are stored at a stack memory location set up or reserved by a calling program, and thus these values are not restored when control is returned to the calling program. As can be seen in pseudo-code 800, the value of register R4 (A0) is stored at memory location M12, and the value of register R5 (A1) is stored at memory location M13.

[0071] After saving the values of the argument registers, the value of return address field 406 is examined. In the case of the example instruction 475A, the encoded value is 1. Thus, the value of the first temporary variable is reduced by four and the value in register R31 (the return address) is saved at memory location M11.

[0072] The next action is to save any additional static registers on stack 1400 as indicated by the encoding in register field 456 of instruction 475A. In an embodiment, the encoding "100" indicates that the values of registers R18, R19, R20, and R21 are to be saved on stack 1400 as static values (e.g., they are to be saved as values associated with storage not yet allocated on the stack). As indicated by pseudo-code 800, the value of register R21 is stored at memory location M10. The values of registers R20, R19, and R18 are stored at memory locations M9, M8, and M7, respectfully.

[0073] Next, the value of static field 410 is examined. In the case of example instruction 475A, the encoded value is 1. Thus, the value of the first temporary variable is reduced by four and the value in register R17 (register S1) is saved at memory location M6.

[0074] The value of static field 408 is examined next. In the case of example instruction 475A, the encoded value is 1. Thus, the value of the first temporary variable is again reduced by four and the value in register R16 (register S0) is saved at memory location M5.

[0075] The next step of pseudo-code 800 is to store any argument registers on stack 1400 that are to be treated as static registers in accordance with the

encoding of field 460. From looking at either FIG. 7 or pseudo-code 800, it can be determined that the values in register R6 (A2) and R7 (A3) are to be saved on stack 1400 as static values. The value of register R7 is saved at memory location M4. The value of register R6 is saved at memory location M3.

[0076] Lastly, the concatenated value of frame-size fields 458 and 412 is examined and used to adjust the value in the stack pointer register, R29 (SP). In the case of example instruction 475A, the encoded value is 20. As indicated by pseudo-code 800, the value in the register R29 is adjusted by shifting the concatenated frame-size value left three bits, padding the new value with zeros, and then subtracting the new value from the value in register R29. At the end of this operation, the stack pointer points to a memory location on stack 1400 below memory location M0.

[0077] FIGs. 15 and 16 illustrate the operation of two additional example save instructions 475B and 475C encoded according to the embodiment of the invention described above. The instructions are assumed to be implemented using a processor operating in accordance with pseudo-code 800. As will be understood by a person skilled in the relevant art given the description herein, the invention allows software tools (e.g., a compiler) to flexibly determine whether values stored in a specified argument register are to be treated as argument values or statics values.

[0078] FIG. 9 illustrates two exemplary 16-bit instructions 902 and 975 that encode method embodiments of the invention for restoring registers and deallocating memory before exit from a subroutine. Instruction 902 can be executed as a single, stand-alone 16-bit instruction. In addition, instruction 902 can be executed in combination with instruction extension 952 as an extended 16-bit instruction or a 32-bit instruction 975.

[0079] Instruction 902 comprises five fields. These five fields are a 9-bit restore instruction opcode field 904, a 1-bit return address register (RA) field 906, a 1-bit first static register (S0) field 908, a 1-bit second static register (S1) field 910, and a 4-bit frame-size field 912. The purpose of each of these fields is described below with regard to the steps of method 1000 and pseudo-code 1100.

[0080] Instruction extension 952 comprises four fields. These four fields are a 5-bit extend instruction opcode field 954, a 3-bit additional static registers field 956, a 4-bit frame-size field 958, and a 4-bit arguments registers field 960. The purpose of each of these fields is described below with regard to the steps of method 1000 and pseudo-code 1200.

[0081] A method 1000 for restoring registers and deallocating memory before exit from a subroutine according to the invention will now be describe. The steps of method 1000 can be encoded by instruction 902 and/or instruction 975 for implementation by a processor.

[0082] FIGs. 10A and 10B illustrate a flowchart of the steps of method 1000. As illustrated in the flowchart, method 1000 comprises a plurality of steps divided up into three stages 1001, 1009, and 1015. Stage 1001 comprises steps 1002, 1004, 1006, and 1008, and is performed prior to or during development and/or implementation of a processor design. As will be understood by a person skilled in the relevant art(s), once these steps have been performed, they need not be repeated. Steps 1002-1008 are an inherent part of a processor, a mapper, a decoder and/or an instruction according to the invention. Stage 1009 comprises steps 1010, 1012, and 1014, and an encoding step implicit in step 1022. These steps of method 1000 are typically performed using a compiler program. Lastly, stage 1015 comprises steps 1016, 1018, 1020, and 1022. These steps are performed by an apparatus or system according to the invention executing an instruction encoded according to the invention. As described herein, it is not necessary to perform each step of method 1000 every time the method is implemented.

[0083] In the description that follows, method 1000 is described with regard to instruction 902, register bank 156, stack 1700 (shown in FIG. 17), and pseudo-code 1100 (shown in FIG. 11). Because many of the steps of method 1000 are similar to the steps of method 500, some of the steps of method 1000 are described with less detail than the steps of method 500.

[0084] Method 1000 begins at a step 1002. In step 1002, at least one register of register bank 156 is specified as an argument register. In an embodiment of the

invention, registers R4, R5, R6, and R7 are specified as argument registers A0, A1, A2, and A3, respectively. In other embodiments, more or less than four argument registers can be specified in step 1002. Because instructions 902 and 975 are intended to be used in conjunction with instructions 402 and 475, the same argument registers should be specified in step 1002 as are specified in step 502 of method 500 above.

[0085] As already described herein, specifying a general purpose register of register bank 156 as an argument register is not intended to alter the nature of the general purpose register with regard to other instructions of an instruction set. A feature of the invention allows any specified argument register according to the invention to be treated as a general purpose register by instructions other than the instructions described herein. General purpose registers thus maintain all the functionality of a general purpose register with regard to other instructions. This feature of the invention applies to general purpose registers specified, for example, as a static register (e.g., S0 and S1), a stack pointer register (SP), and/or a return address register (RA) in accordance with method 1000.

[0086] In step 1004, at least one register of register bank 156 is specified as a static register. In the embodiment of FIG. 2, two registers R16 and R17 are specified as static registers S0 and S1, respectively. In other embodiments, more or less than two registers may be specified. Again, because instructions 902 and 975 are intended to be used in conjunction with instructions 402 and 475, the same static registers should be specified in step 1004 as are specified in step 504 of method 500 above.

[0087] In step 1006, a register of register bank 156 is specified as a stack pointer register. In the embodiment shown in FIG. 2, register R29 of register bank 156 is specified as the stack pointer register.

[0088] In step 1008, a register of register bank 156 is specified as a return address register. In the embodiment shown in FIG. 2, register R31 is specified as the return address register.

[0089] In step 1010, fields 908 and 910 of instruction 902 are encoded to indicate whether static registers R16 and R17 are to be restored from stack 1700 before

[0090] Step 1012 of method 1000 is not performed by the method embodiment encoded by instruction 902. For the method embodiment encoded by instruction 902, it is assumed that each of the registers specified in step 1002 is to be treated as an argument register. This is not the case, however, for other method embodiments of the invention encoded, for example, using instruction 975. Step 1012 is further described below with regard to instruction 975.

[0092] Step 1016 of method 1000 is also not performed by the method embodiment encoded by instruction 902. As described with regard to step 1012, it is assumed that each of the registers specified in step 1002 is to be treated as an argument register. Thus, none of the argument registers R4, R5, R6, or R7 is restored before exit from a subroutine.

[0093] In step 1018, a static value is either loaded or not loaded into register R16 from stack 1700 when instruction 902 is executed. Whether a static value is loaded into register R16 is based on the encoding of field 908. In an embodiment, the static value of register R16 is restored only if bit 5 of instruction 902 is set. Otherwise, the value of register R16 will not be restored when instruction 902 is executed. Similarly, in step 1018, a static value of register R17

will be restored from stack 1700 when instruction 902 is executed only if bit 4 of instruction 902 is set. Otherwise, the value of register R17 will not be restored when instruction 902 is executed.


[0094] In step 1020, a return address value is either loaded or not loaded into register R31 from stack 1700 when instruction 902 is executed. Whether the return address value is loaded into register R31 is based on the encoding of field 906. In an embodiment, the return address value is loaded into register R31 only if bit 6 of instruction 902 is set. Otherwise, the return address value is not loaded into register R31 when instruction 902 is executed.

[0095] Lastly, in step 1022, the value in the stack pointer register (R29) is adjusted based on a value encoded in field 912 of instruction 902. Adjusting the stack pointer allows a stack frame to be deallocated before exit from a subroutine. In an embodiment, the value in the stack pointer register is adjusted by adding 8 times the value encoded in field 912 of instruction 902 to the value of register R29. If the value encoded in field 912 is zero, the value in the stack pointer register is adjusted by adding 128. In other embodiments, the value in the stack pointer register is adjusted by different amounts.

[0096] FIG. 11 illustrates pseudo-code 1100 according to the invention. Pseudo-code 1100 describes to a person skilled in the relevant art how a processor according to the invention implements the steps of method 1100 encoded in an instruction 902.

[0097] The steps of pseudo-code 1100 will now be described. The steps of pseudo-code 1100 are described with reference to register bank 156, stack 1700, and example instruction 902A shown below and in FIG. 17. Example instruction 902A can be used to restore register values save to the stack by example instruction 402A, described above. As can be seen by examining example instruction 902A, the encodings of the instruction fields are: "1" for return address field 906; "1" for static register field 908; "1" for static register field 910; and "0100" (i.e., 4) for frame-size field 912. The encoding for the restore instruction opcode field 904 is shown as "XXXXXXXXXX" to indicate that the

encoding is a processor specific value that is unimportant for understanding the steps of pseudo-code 1100.



[0098] Implementation of pseudo-code 1100 starts by adjusting the value of register R29 (the stack pointer) and saving the new value to a temporary variable. At the start of pseudo-code 1100, the stack pointer is assumed to point to memory location M4. The new value saved in the temporary variable points to memory location M12.

[0099] After storing the adjusted stack pointer value to a temporary variable, the value of return address field 906 is examined. In the case of example instruction 902A, the encoded value is 1. Thus, the value of the temporary variable is reduced by four and the value at memory location M11 (the return address) is loaded into register R31.

[0100] Next, the value of static field 910 is examined. In the case of example instruction 902A above, the encoded value is 1. Thus, the value of the temporary variable is again reduced by four and the value saved at memory location M10 is loaded into register R17 (register S1). Similarly, the value of static field 908 is examined. In the case of the example instruction 902A, the encoded value is 1. Thus, the value of the temporary variable is again reduced by four and the value saved at memory location M9 is loaded into register R16 (register S0).

[0101] Lastly, as indicated by pseudo-code 1100, the value in the stack pointer register (R29) is adjusted to deallocate the stack frame (e.g., the stack frame created using example instruction 402A).

[0102] Method 1100 will now be described with regard to instruction 975, register bank 156, stack 1800 (shown in FIG. 18), and pseudo-code 1200 (shown in FIGs. 12A and 12B) in order to point out differences between instruction 975 and instruction 902.

[0103] As will be understood by a person skilled in the relevant art given the description herein, steps 1002, 1004, 1006, and 1008 do not change when

instruction 975 is used to encode method 1000 rather than instruction 902. Thus, the description for these steps is not repeated here.

[0104] In step 1010, fields 908 and 910 of instruction 975 are encoded to indicate whether static registers R16 and R17 are to be restored before exit from a subroutine. In an embodiment, if bits 4 and 5 of instruction 975 are set, the static values on stack 1800 will be restored to registers R16 and R17. For example, if bit 4 of instruction 902 is set, a value at memory location M5 will be loaded into register R16. Similarly, if bit 5 of instruction 902 is set, a value at memory location M6 will be loaded into register R17.

[0105] In step 1012 of method 1000, a 4-bit binary value is encoded in argument registers field 960 (*aregs*) of instruction extension 952 to indicate which registers specified as argument registers in step 902 are to be restored as static registers during execution of instruction 975. FIG. 7 illustrates the 4-bit binary encoding used in an embodiment of the invention. Other encodings, however, can also be used in accordance with method 1000.

[0106] In step 1014, field 906 of instruction 975 is encoded to indicate whether a value on stack 1800 is to be loaded into return address register R31 before exit from a subroutine. In an embodiment, for example, a value stored at memory location M11 will be loaded into register R31 during execution of instruction 975 if bit 6 of instruction 975 is set. If bit 6 is cleared, no value is loaded into register R31 during execution of instruction 975.

[0107] In step 1016, none, one, two, three, or four of the argument registers R4, R5, R6, or R7 specified in step 1002 are restored during execution of instruction 975. Which, if any, argument registers are restored depends on the value encoded in field 960 of instruction 975. FIG. 7 shows which specified argument registers are restored as static registers for the 4-bit binary encoding illustrated in FIG. 7.

[0108] In step 1018, register R16 is either restored or not restored during execution of instruction 975. Whether register R16 is restored is based on the encoding of field 908. In an embodiment, register R16 is restored using a value from stack 1800 only if bit 5 of instruction 975 is set. Otherwise, register R16 is not restored when instruction 975 is executed. Similarly, in step 1018, register

R17 is either restored or not restored during execution of instruction 975 based on the encoding of field 910. In an embodiment, register R17 is restored using a value from stack 1800 only if bit 4 of instruction 975 is set. If bit 4 is zero, register R17 is not restored when instruction 975 is executed.

[0109] In step 1020, register R31 is either restored or not restored during execution of instruction 975 based on the encoding of field 906. In an embodiment, register R31 is restored using a value from memory location M11 of stack 1800 only if bit 6 of instruction 975 is set. Otherwise, register R31 is not restored when instruction 975 is executed.



[0110] In step 1022, the value in register R29 is adjusted based on a value encoded in fields 958 and 912 of instruction 975. Adjusting the stack pointer value in register R29 allows stack memory to be deallocated before exit from a called subroutine. In an embodiment, the value in register R29 is adjusted by adding 8 times the frame-size value encoded in fields 958 and 912 of instruction 975. The 4-bits of field 958 and the 4-bits of field 912 are concatenated to form an 8-bit frame-size value. In other embodiments, the value in register R29 is adjusted by different amounts.

[0111] Field 956 of instruction 975 is used to encode whether additional registers of register bank 156 are restored as static registers before exit from a subroutine. For example, in an embodiment, registers R18, R19, R20, R21, R22, R23, and R30 can be restored as additional static registers. This feature of the invention is further described below with reference to pseudo-code 1200.

[0112] FIGs. 12A and 12B illustrate pseudo-code 1200 according to the invention. Pseudo-code 1200 describes to a person skilled in the relevant art how a processor according to the invention implements the steps of method 1000 encoded in an instruction 975.

[0113] The steps of pseudo-code 1200 will now be described. The steps of pseudo-code 1200 are described with reference to register bank 156, stack 1800, and example instruction 975A shown below and in FIG. 18. Example instruction 975A can be used to restore static values saved to the stack using example instruction 475A. As can be seen by examining the example instruction, the

encodings of the instruction fields are: "1" for return address field 906; "1" for static register field 908; "1" for static register field 910; "00010100" (i.e., 20) for the concatenated frame-size fields 958 and 912; "100" for the additional static registers field 956; and "1010" for the argument registers field 960. The encodings for the restore instruction opcode field 904 and the extend instruction opcode field 954 are shown as "XXXXXXXXXX" and "XXXXXX," respectfully, to indicate that these encodings are processor specific values that are unimportant for understanding the steps of pseudo-code 1200.



[0114] Implementation of pseudo-code 1200 starts by adjusting the value of register R29 (the stack pointer), based on a frame-size value encoded in fields 958 and 912, and saving the new value to two temporary variables. As indicated by pseudo-code 1200, the value in the stack pointer register (R29) is adjusted by shifting the concatenated frame-size value left three bits, padding the new value with zeros, and then adding the new value to the value in register R29. At the start of pseudo-code 1200, the stack pointer is assumed to be pointing at a memory location below memory location M0 (i.e, 64 bytes below memory location M0). The new value saved in the temporary variables points to memory location M12.

[0115] After storing the adjusted stack pointer value to a first and second temporary variable, the value of return address field 906 is examined. In the case of example instruction 975A, the encoded value is 1. Thus, the value of the first temporary variable is reduced by four and the value at memory location M11 (the return address) is loaded into register R31.

[0116] After restoring the return address, pseudo-code 1200 restores any extra static registers encoded in field 956 of instruction 975A. As described herein, the encoding "100" indicates that registers R21, R20, R19, and R18 are to be restored as extra static register. Thus, in an embodiment, the values stored at memory locations M10, M9, M8, and M7 are loaded into registers R21, R20, R19, and R18, respectfully.

[0117] Next, the value of static field 910 is examined. In the case of example instruction 975A, the encoded value is 1. Thus, the value of the first temporary variable is again reduced by four and the value saved at memory location M6 is loaded into register R17 (register S1). Similarly, the value of static field 908 is examined. In the case of example instruction 975A, the encoded value is 1. Thus, the value of the first temporary variable is again reduced by four and the value saved at memory location M5 is loaded into register R16 (register S0).

[0118] The next step of pseudo-code 1200 is to restore any argument register values saved on stack 1800 that are to be treated as static values in accordance with the encoding of field 960. From looking at either FIG. 7 or pseudo-code 1200, it can be determined that registers R6 and R7 are to be restored as statics registers. Accordingly, register R7 is restored with the value saved at memory location M4. Register R6 is restored with the value saved at memory location M3.

[0119] Lastly, as indicated by pseudo-code 1200, the value in the stack pointer register (R29) is adjusted by loading register R29 with the adjusted stack pointer value saved in the second temporary variable. It should be noted that registers R4 and R5 are not restored using the values stored at memory locations M13 and M12.

[0120] Based on the description of the invention herein, a person skilled in the relevant art will understand how to implement the invention to flexibly allocate which values contained in the registers of register bank 156 are saved as static or argument values upon entry to a subroutine, and which registers are restored as static registers before exit from a subroutine. A person skilled in the relevant art

will understand based on the description herein how to implement the invention to flexibly allocate and deallocate stack memory according to the invention.

CONCLUSION

[0121] While various embodiments of the present invention have been described above, it should be understood that they have been presented by way of example, and not limitation. It will be apparent to persons skilled in the relevant art(s) that various changes in form and detail can be made therein without departing from the spirit and scope of the invention.

[0122] In addition to implementations using hardware (e.g., within a microprocessor or microcontroller), implementations also may be embodied in software disposed, for example, in a computer usable (e.g., readable) medium configured to store the software (i.e., a computer readable program code). The program code causes the enablement of the functions or fabrication, or both, of the systems, techniques and operations disclosed herein. For example, this can be accomplished through the use of general programming languages (e.g., C or C++), hardware description languages (HDL) including Verilog HDL, VHDL, and so on, or other available programming and/or circuit (i.e., schematic) capture tools. The program code can be disposed in any known computer usable medium including semiconductor, magnetic disk, optical disk (e.g., CD-ROM, DVD-ROM) and as a computer data signal embodied in a computer usable (e.g., readable) transmission medium (e.g., carrier wave or any other medium including digital, optical, or analog-based medium). As such, the code can be transmitted over communication networks including the Internet and intranets.

[0123] It is understood that the functions accomplished and/or structure provided by the embodiments and techniques described above can be represented in a core (e.g., a microprocessor core) that is embodied in program code and may be transformed to hardware as part of the production of integrated circuits. Also, the system, techniques, and functionality may be embodied as a combination of hardware and software. Moreover, the save and restore functionality described

[illegible]